

Model-based optimization of MPDATA on Intel Xeon Phi through load imbalancing

Alexey Lastovetsky
University College Dublin
Belfield, Dublin 4, Ireland
alexey.lastovetsky@ucd.ie

Lukasz Szustak
Czestochowa University of
Technology
Dabrowskiego 69, 42-201
Czestochowa, Poland
lszustak@icis.pcz.pl

Roman Wyrzykowski
Czestochowa University of
Technology
Dabrowskiego 69, 42-201
Czestochowa, Poland
roman@icis.pcz.pl

ABSTRACT

Load balancing is a widely accepted technique for performance optimization of scientific applications on parallel architectures. Indeed, balanced applications do not waste processor cycles on waiting at points of synchronization and data exchange, maximizing this way the utilization of processors. In this paper, we challenge the universality of the load-balancing approach to optimization of the performance of parallel applications. First, we formulate conditions that should be satisfied by the performance profile of an application in order for the application to achieve its best performance via load balancing. Then we use a real-life scientific application, MPDATA, to demonstrate that its performance profile on a modern parallel architecture, Intel Xeon Phi, significantly deviates from these conditions. Based on this observation, we propose a method of performance optimization of scientific applications through load imbalancing. We also propose an algorithm that finds the optimal, possibly imbalanced, configuration of a data parallel application on a set of homogeneous processors. This algorithm uses functional performance models of the application to find the partitioning that minimizes its computation time but not necessarily balances the load of the processors. We show how to apply this algorithm to optimization of MPDATA on Intel Xeon Phi. Experimental results demonstrate that the performance of this carefully optimized load-balanced application can be further improved by 15% using the proposed load-imbalancing optimization.

Keywords

functional performance model, data partitioning, Intel Xeon Phi, MPDATA, load imbalancing

1. INTRODUCTION

Load balancing is a widely accepted technique for optimization of the computation performance of scientific applications on parallel architectures. Indeed, the intuition suggests that unlike unbalanced applications, the balanced ones

do not waste processor cycles on waiting at points of synchronization and data exchange, maximizing this way the utilization of the processors.

In this paper, we challenge the universality of the load-balancing approach to optimization of the computation performance of parallel applications. First, we try to understand the limitations of the load-balancing approach. We formulate conditions that should be satisfied by the performance profile of an application in order for the application to achieve its best performance via load balancing.

Then we use a real-life scientific application to demonstrate that its performance profile on a modern parallel architecture does not satisfy these conditions. The application we use implements the Multidimensional Positive Definite Advection Transport Algorithm (MPDATA), which is one of the major parts of the dynamic core of the EULAG geophysical model [21]. EULAG (Eulerian/semi-Lagrangian fluid solver) is an established numerical model developed for simulating thermo-fluid flows across a wide range of scales and physical scenarios [19, 22]. In particular, it can be used in numerical weather prediction (NWP), simulation of urban flows, areas of turbulence, ocean currents, etc. This solver, originally developed for conventional HPC systems, is currently being re-written for modern HPC platforms. In particular, MPDATA has been recently re-written and optimized for execution on an Intel Xeon Phi coprocessor [25, 26].

In our experiments, we observe significant deviations of the MPDATA performance profile from the conditions required for applicability of the load-balancing techniques. Based on this observation, we propose a general method of performance optimization of scientific applications through load imbalancing as well as an algorithm that finds the optimal, possibly imbalanced, configuration of a data parallel application on a set of homogeneous processors. This algorithm uses functional performance models of the application [14, 13] to find the partitioning that minimizes its computation time but not necessarily balances the load of the processors. Finally, we apply this algorithm to optimization of MPDATA on Intel Xeon Phi. Experimental results demonstrate that the performance of this carefully optimized load-balanced application can be further improved by 15% using the proposed load-imbalancing method.

The contributions of the work presented in the paper are as follows:

- Formulation of the conditions that should be satisfied to guarantee that load balancing will minimize the computation time of parallel application.
- Building the performance profile of a real-life scientific application on a modern HPC platform and demonstration of its significant deviation from the conditions that guarantee that load balancing be a safe technique for performance optimization.
- A new optimization method that uses the performance profile for optimization of the application through its imbalancing.
- A partitioning algorithm finding the optimal and generally speaking uneven distribution of computations of an application between homogeneous processing units using its functional performance model.
- Application of the proposed partitioning algorithm to optimization of MPDATA on Intel Xeon Phi, resulting in further acceleration of this carefully optimized load-balanced application by up to 15%.

The rest of the paper is structured as follows. Section 2 overviews load-balancing techniques and formulates the conditions when these techniques would minimize the computation time of parallel applications. Section 3 analyzes the performance profile of MPDATA on Intel Xeon Phi and introduces the new approach to minimization of the computation time through load imbalancing. Section 4 introduces a partitioning algorithm for (uneven) distribution of computations between homogeneous processors, minimizing the computation time of the application. Section 5 applies this algorithm to optimization of MPDATA on Xeon Phi. Section 6 presents experimental results, and Section 7 concludes the paper.

2. LOAD BALANCING AND PERFORMANCE

In this section we overview load-balancing techniques used for optimization of the performance of parallel scientific applications on both homogeneous and heterogeneous platforms. We also formulate the conditions when application of these techniques will optimize the computation performance.

Load balancing is a widely accepted and commonly used approach to performance optimization of scientific applications on parallel architectures. Load balancing algorithms can be classified as static or dynamic. Static algorithms (for example, those based on data partitioning) [8, 27, 18, 13] require a priori information about the parallel application and platform. This information can be gathered either at compile-time or runtime. Static algorithms are also known as predicting-the-future because they rely on accurate performance models as input to predict the future execution of the application. Static algorithms are particularly useful for applications where data locality is important because they do not require data redistribution. However, these algorithms are unable to balance on non-dedicated platforms, where load changes with time. Dynamic algorithms (such

as task scheduling and work stealing) [15, 3, 20] balance the load by moving fine-grained tasks between processors during the calculation. Dynamic algorithms do not require a priori information about execution but may incur significant communication overhead due to data migration. Dynamic algorithms often use static partitioning for their initial step due to its provably near-optimal communication cost, bounded tiny load imbalance, and lesser scheduling overhead [24].

Whatever load balancing algorithm is used, the goal is always to minimize the computation time of the application. The intuition behind the assumption that balancing the application improves its performance is the following: a balanced application does not waste processor cycles on waiting at points of synchronization and data exchange, maximizing this way the utilization of the processors. Is this assumption always true? To answer this question, let us formulate the assumption in a mathematical form. Consider an application, the computation performance of which can be modeled by speed functions. Namely, let p parallel processors be used to execute the application, and let $s_i(x)$ be the speed of execution of the workload of size x by processor i . Here the speed can be measured in floating point operations per second or any other fix-sized computation units per unit time. The size of workload can be characterized by the problem size (say, the number of cells in the computational domain or the matrix size) or just by the number of equal-sized computational units. Anyway, the speed $s_i(x)$ is calculated as $\frac{x}{t_i(x)}$, where $t_i(x)$ is the execution time of the workload of size x on processor i . Using these definitions, we can formulate the following theorem.

Theorem 1: Let $s_i(x) > 0$ ($x > 0$) be the speed of processor $i \in \{1, \dots, p\}$, and $\forall \Delta x > 0$: $\frac{s_i(x)}{x} \geq \frac{s_i(x+\Delta x)}{x+\Delta x}$. Let $x_1 + \dots + x_p = n > 0$ and $\frac{s_1(x_1)}{x_1} = \dots = \frac{s_p(x_p)}{x_p}$. Then, $\forall y_1, \dots, y_p > 0$ such that $(y_1, \dots, y_p) \neq (x_1, \dots, x_p)$ and $y_1 + \dots + y_p = n$: $\max_i \frac{y_i}{s_i(y_i)} \geq \frac{x_1}{s_1(x_1)}$.

Proof: As $(y_1, \dots, y_p) \neq (x_1, \dots, x_p)$ and $y_1 + \dots + y_p = x_1 + \dots + x_p$, then there exists $k \in \{1, \dots, p\}$ such that $y_k > x_k$. Therefore, $\max_i \frac{y_i}{s_i(y_i)} \geq \frac{y_k}{s_k(y_k)} = \frac{x_k + (y_k - x_k)}{s_k(x_k + (y_k - x_k))} \geq \frac{x_k}{s_k(x_k)} = \frac{x_1}{s_1(x_1)}$.

Theorem 1 states that in order to guarantee that the balanced configuration of the application will execute the workload of size n faster than any unbalanced configuration, the speed functions $s_i(x)$, characterizing the performance profiles of the processors, should satisfy the condition:

$$\forall \Delta x > 0: \frac{s_i(x)}{x} \geq \frac{s_i(x + \Delta x)}{x + \Delta x} \quad (1)$$

Geometrically, it can be illustrated as follows. If we plot a speed function as shown in Figure 1, then the angle $\alpha(x)$ between the straight line, connecting the point $(0,0)$ and the point $(x, s(x))$ on the speed curve, and the x -axis will be inversely proportional to the execution time of the workload of size x by the processor. Indeed, the cotangent of this angle is directly proportional to the ratio $\frac{x}{s(x)}$ representing the execution time of the workload x . Therefore, larger angles correspond to shorter execution times. The condition 1

means that the increase of the workload, x , will never result in the decrease of the execution time, or equivalently in the increase of the angle $\alpha(x)$.




Figure 1: Example of speed function suitable for minimization of computation time through load balancing. Angle $\alpha(x)$ represents the computation time: the greater the angle, the shorter the computation time.

The main body of the load balancing algorithms designed for performance optimization explicitly or implicitly assume that the speed of processor does not depend on the size of workload [8, 27, 16, 4, 10, 17]. In other words, the speed functions $s_i(x)$ are assumed to be positive constants, in which case the condition 1 is trivially satisfied. More advanced algorithms are based on functional performance models (FPMs), which represent the speed of processor by a continuous function of the problem size [14, 11]. However, the shape of the function is not arbitrary but has to satisfy the following assumption [13]: Along each of the problem size variables, either the function is monotonically decreasing, or there exists point x such that

- On the interval $[0, x]$, the function is
 - monotonically increasing,
 - concave, and
 - any straight line coming through the origin of the coordinate system intersects the graph of the function in no more than one point.
- On the interval $[x, \infty)$, the function is monotonically decreasing.

These restrictions on the shape of speed functions guarantee that the efficient load balancing algorithms, proposed in [9, 7, 12, 5, 6, 2], will always return a unique solution, minimizing the computation time. At the same time, it is easy to show that the restrictions imposed on FPMs will make them comfortably satisfy the condition 1.

Thus, the state-of-the-art load balancing algorithms designed for optimization of the computation performance of parallel applications assume that their performance profiles satisfy the condition 1. Therefore, correct application of such algorithms requires that the experimental speed points be approximated by a function satisfying this condition. This approximation step may significantly distort the actual performance profile and lead to a substantially non-optimal solution.

3. OPTIMIZATION OF PARALLEL APPLICATIONS THROUGH LOAD IMBALANCING

In this section, we demonstrate that the performance profile of real-life scientific applications on modern parallel platforms may significantly deviate from the conditions, which guarantee that load balancing will always optimize their computation performance. Based on this observation, we propose an optimization method that uses the performance profile for optimization of the application through its imbalancing.

In this work, we build the performance profile of MPDATA on Intel Xeon Phi. MPDATA is a core component of EULAG (Eulerian/semi-Lagrangian fluid solver), which is an established computational model developed for simulating thermo-fluid flows across a wide range of scales and physical scenarios. Its carefully optimized data-parallel implementation on a 60-core Intel Xeon Phi [26] partitions the 3D rectilinear $n \times n \times l$ domain into four equal $\frac{n}{2} \times \frac{n}{2} \times l$ sub-domains, each allocated to a team of 15 cores. This configuration of the application is the best load-balanced configuration identified in [26]. The experimentally constructed speed functions of these four teams, each processing (in parallel) a $120 \times m \times 128$ sub-domain, are shown in Figure 2.




Figure 2: Speed functions of Intel MIC built for four 15-core teams, each processing in parallel a $120 \times m \times 128$ sub-domain. The speed is measured in cells per second, while the problem size is represented by m .

This graph clearly shows that for many m and Δm the speed of processing of the $120 \times m \times 128$ sub-domain will be significantly lower than the speed of processing of the $120 \times (m + \Delta m) \times 128$ sub-domain. Moreover, we can also see that $\alpha(m + \Delta m) > \alpha(m)$ for some such m and Δm , which means that the time of processing of the $120 \times m \times 128$ sub-domain will be longer than the time of processing of the $120 \times (m + \Delta m) \times 128$ sub-domain. The latter observation can be used to speed up the execution of the application. Namely, if we re-partition the equally partitioned domain so that two teams get $120 \times (m + \Delta m) \times 128$ sub-domains and two other teams get $120 \times (m - \Delta m) \times 128$ sub-domains, and $\min\{\alpha(m + \Delta m), \alpha(m - \Delta m)\} > \alpha(m)$, then this unequal and unbalanced partitioning will result in faster execution.

In general, if the performance profile of an application vio-

lates the condition 1, that is,

$$\exists i \in \{1, \dots, p\} \exists x > 0 \exists \Delta x > 0: \frac{s_i(x)}{x} < \frac{s_i(x + \Delta x)}{x + \Delta x} \quad (2)$$

and the balanced configuration of the application allocates the workload of size x to processor i , then the application can be accelerated if we reduce the accumulated workload of all processors but processor i by Δx so that none of these processors would increase its execution time, and allocate this additional workload to processor i . This method can be applied to optimization of parallel applications on both heterogeneous and homogeneous platforms.

4. MODEL-BASED PARTITIONING ALGORITHM FOR OPTIMAL LOAD IMBALANCING

In this section, we develop the proposed approach for a relatively simple case and introduce a partitioning algorithm that finds the optimal distribution of computations of an application between homogeneous processors using the functional performance model of the application.

Consider the following problem. Let p identical parallel processors be used to execute the workload of size n , and let $s(x)$ be the speed of execution of the workload of size x by a processor. Let Δx be the minimal granularity of workload so that each processor can be only allocated a multiple of Δx . The problem is to find the distribution of the workload of size n between the p processors, which minimizes the computation time of its parallel execution.

For simplicity, we assume that $\frac{n}{p}$ be a multiple of Δx and p be an even number. Then the following Algorithm 1 will solve this problem.

Algorithm 1 Optimal distribution of workload between homogeneous processors

```

 $x_{ropt} = x_{lopt} = x_r = x_l = \frac{n}{p}$ 
 $t_{min} = \frac{\frac{n}{p}}{s(\frac{n}{p})}$ 
repeat
   $x_r += \Delta x$ 
   $x_l -= \Delta x$ 
   $t = \max(\frac{x_r}{s(x_r)}, \frac{x_l}{s(x_l)})$ 
  if  $t < t_{min}$  then
     $t_{min} = t$ 
     $x_{ropt} = x_r$ 
     $x_{lopt} = x_l$ 
  end if
until  $x_r < n$ 

```

This algorithm returns the optimal distribution of the workload of size n between p processors so that each odd processor receives the workload of size x_{lopt} , and each even processor receives the workload of size x_{ropt} . It can be proved that the solution returned by this algorithm will always minimize the execution time of the given workload n . Note that the traditional load-balanced approach would assign the workload of size $\frac{n}{p}$ to all processors.

It is obvious that if we replace the speed function $s(x)$ by any function $a \times s(x)$, where $a = const$, then this algorithm

will return the same solution. We will use this property when applying Algorithm 1 in Section 5.4.

Algorithm 1 can be easily generalized for an arbitrary (not only even) number of processors.

5. APPLICATION

In this section, we apply the partitioning algorithm proposed in Section 4 to optimization of MPDATA on Intel Xeon Phi.

5.1 Intel MIC overview

The Intel MIC architecture is a relatively new system for high performance computing [1]. Intel MIC combines many integrated Intel CPU cores into a single chip. This architecture is built to provide a general-purpose programming environment similar to that provided for Intel CPUs. It is capable of running applications written in industry-standard programming languages such as Fortran, C, and C++. The Intel Xeon Phi (codenamed Knights Corner) is the first product based on Intel MIC architecture. This coprocessor is delivered in form factor of a PCI express device, and can not be used as a stand-alone processor. However, it allows users to directly run individual applications in the native mode without the support of CPU.

In this study, we use the top-of-the-line Intel Xeon Phi 7120P coprocessor. It contains 61 cores clocked at 1.238 GHz, and 16 GB of on-board memory. As the Intel MIC architecture supports four-way hyper-threading, it totally gives 244 logical cores (threads) for a single chip. This coprocessor provides 352 GB/s of memory bandwidth. An important component of each Intel Xeon Phi processing core is its vector processing unit (VPU) [26], that significantly increases the computing power. Each VPU supports a new 512-bit SIMD instruction set called Intel Initial ManyCore Instructions. The theoretical peak performance of Intel Xeon Phi 7120P is 1208 GFlop/s for double precision numbers.

5.2 Introduction to MPDATA

The MPDATA algorithm is a general approach for integrating the conservation laws of geophysical fluids on micro-to-planetary scales [23]. It belongs to the class of methods for the numerical simulation of fluid flows which are based on the sign-preserving properties of upstream differencing. The MPDATA scheme allows for solving advection problems, and offers several options to model a wide range of complex geophysical flows.

MPDATA corresponds to the group of nonoscillatory forward-in-time algorithms. The number of required time steps depends on a type of simulated physical phenomenon, and can exceed few millions especially when considering the MPDATA algorithm as a part of the EULAG model. For detailed description of the MPDATA mathematical scheme, the reader is referred to [21, 22, 23].

Each MPDATA time step is determined by a set of 17 computational stages, where each stage is responsible for calculating elements of a certain matrix. These stages represent stencil codes which update grid elements according to different patterns. Listing 1 shows a part of the 3D MPDATA stencil-based implementation for the 8-th stage.

Listing 1: Part of 3D MPDATA stencil-based implementation

```

/*...*/
//stage 8
for( ... ) // i - dimension
    for( ... ) // j - dimension
        for( ... ) // k - dimension
            mx[i,j,k]=max(x[i][j][k],
                           x[i-1][j][k], x[i+1][j][k],
                           x[i][j-1][k], x[i][j+1][k],
                           x[i][j][k-1], x[i][j][k+1]);
/*...*/

```

The stages are dependent on each other: outcomes of prior stages are usually input data for the subsequent computations. Every stage reads a required set of matrices from the main memory, and writes results to the main memory after computation. In consequence, a significant traffic to the main memory is generated, which mostly limits the performance of novel architectures. A single MPDATA time step requires 5 input matrices, and returns one output matrix that is necessary for the next step.

5.3 Adaptation of 3D MPDATA to Intel Xeon Phi coprocessor

In our previous work [25, 26], we proposed the adaptation of 3D MPDATA to Intel Xeon Phi coprocessors. The proposed decomposition contributes to ease the memory and communication bounds, and to better exploit computation resources of Intel Xeon Phi. The resulting adaptation is based on the following methodology:

- (3+1)D decomposition of MPDATA heterogeneous stencil computations;
- partitioning of threads into independent work teams;
- parallelization of MPDATA computations;
- scheduling for multicore and manycore systems.

To workaroud the memory-bound nature of MPDATA, we proposed the (3+1)D decomposition of MPDATA stencil computation [26]. The main aim of this decomposition is to take advantage of cache memory reuse by transferring the data traffic associated with all intermediate computation from the main memory to the cache hierarchy. This aim is achieved by using a combination of two well-known loop optimization techniques: loop tiling and loop fusion. Such an approach allows us to reduce the main memory traffic at the cost of additional computations associated with extra areas (*halo*) of all intermediate matrices. Another advantage of this approach is the possibility of reducing the main memory consumption because all intermediate results are stored only in the cache memory.

The proposed decomposition contributes to the data traffic from the main memory to the cache hierarchy. In consequence, a lot of inter- and intra-cache communications are generated between more than 200 Intel MIC's processing cores. To improve the efficiency of the (3+1)D decomposition on Intel Xeon Phi, we provided partitioning of available

cores (threads) into independent work teams. As a result, the MPDATA computing domain is partitioned into P sub-domains of different sizes, each of which is processed by a single work team of threads, according to the proposed (3+1)D decomposition. The number of cores (threads) assigned to different teams can also be different. Figure 3 shows an example of partitioning of 60 Intel MIC's processing cores into 4 teams, and partitioning of the MPDATA grid into 4 pieces as well.




Figure 3: Partitioning of Intel MIC's processing cores into 4 work teams

Within every time step, the work teams execute computations in parallel and independently of each other. After each time step, the work teams are synchronized. Each sub-domain is further partitioned into a number of MPDATA blocks, where subsequent blocks are processed one by one, and each block is processed in parallel by the corresponding work team. Every block is further decomposed into sub-blocks, where each sub-block is processed by a certain thread of the work team. A sequence of all the MPDATA stages is executed within every sub-block, taking into account the data dependencies. This is achieved at the cost of some extra computations performed for halo regions by all teams.

Our best performance results on a single Intel Xeon Phi have been achieved so far by partitioning the 3D $n \times m \times l$ MPDATA domain in two dimensions n and m into four equal sub-domains, so that there is one-to-one mapping between these sub-domains and the teams of cores arranged in a 2×2 grid as illustrated in Figure 3. This partitioning allowed us to balance the load of the core teams and minimize the execution time of the application in comparison with all other partitioning shapes.

5.4 Applying model-based partitioning algorithm to MPDATA decomposition

Thus, the best load-balanced configuration of the MPDATA application on a Intel Xeon Phi arranges its cores in four 15-core teams as shown in Figure 3 and evenly partitions the $n \times m \times l$ computation domain between these teams, allocating a $\frac{n}{2} \times \frac{m}{2} \times l$ sub-domain to each of the teams. In this subsection, we use the data-partitioning algorithm, presented in Section 4, to find a better partitioning of the computation domain between these teams of cores.

As a first step, we build speed functions of the teams so that the speed of each team be represented by a function of problem size. In the case of MPDATA, the problem size is characterized by the size of the domain processed by the team and therefore represented by three parameters n , m and l .

In real-life NWP simulations l is fixed [26]. Therefore, we build speeds of teams as functions of two parameters n and m , setting l to 128, the value typically used in simulations.

In general, the speed should be measured in equal-sized computation units performed per one time unit [13], for example, in flops. In the case of MPDATA, it is difficult to estimate the amount of arithmetic operations that will be executed during the processing of a $n \times m \times l$ computation domain. We know however that with a very high level of accuracy this amount is directly proportional to the number of cells in this domain. Therefore, we measure the speed in cells per second.

The speed functions are built empirically by benchmarking the work teams for a range of problem sizes. For each problem size (n, m) , the speed is calculated as $\frac{n \times m \times 128}{t}$, where t is the measured execution time.

It has been shown [29, 28] that in modern multicore, many-core and hybrid platforms, where processing elements are coupled and share resources, the speed of one group of elements may depend on the load of others due to resource contention. Therefore, the groups cannot be considered as independent processing units and their speed cannot be measured separately and independently. In this work, we use the performance measurement method proposed in [28]. According to this method, the performance of the four teams of cores is measured simultaneously rather than separately, thereby taking into account resource contention. To ensure the reliability of measurements, we repeat measurements multiple times. We only measure the computation time of every team without the overheads of inter-team synchronization required after each time step. If the measurements were conducted separately, the measured performance of these teams would not reflect their actual performance during the execution of the application, and therefore performance optimization decisions based on the corresponding performance models would be inaccurate. Figure 4 demonstrates the difference between the speed of team T_0 measured separately and simultaneously with other teams.

Figure 5 illustrates the speed of team T_0 as a function of parameters n and m (remember that $l = 128$). The experimental points for the speed function were obtained with steps $\Delta n = \Delta m = 4$ for both n and m parameters.

We can see that for a fixed value of m the speed varies very slowly and very little with variation of n , staying nearly constant. More detailed analysis of the speed functions confirms that the speed of team strongly depends on m and very little depends on n . This observation allows us to assume that with a high level of accuracy the optimal (or at least a near optimal) partitioning of a $N \times M \times 128$ domain between the four teams can be obtained from the optimal even load-balanced partitioning, which allocates a sub-domain of size $\frac{N}{2} \times \frac{M}{2} \times 128$ to each team, by fixing the first parameter n to $\frac{N}{2}$ and varying m .

Mathematically, it means that we only have to deal with speed functions of just one parameter, m . These functions are obtained from the previously built speed functions of two parameters, n and m , by fixing the parameter n . Geometri-




Figure 4: Comparison of speed functions of team T_0 , measured separately ($S_{T_0}^*(x)$) and simultaneously with other three teams ($S_{T_0}(x)$) executing the same workload




Figure 5: Experimentally built speed of execution of the MPDATA workload by team T_0 as function of two parameters n and m ($l = 128$)

cally, this can be illustrated as follows. The two-parameter speed functions are represented by surfaces. By fixing parameter n to $\frac{N}{2}$, we cut the surfaces by a vertical plane $n = \frac{N}{2}$ as shown in Figure 5, obtaining on this plane curves, representing the one-parameter speed functions as shown in Figure 6 for $n = 120$.

Finally, as all four teams have very close speed functions (as can be seen in Figure 6), we calculate their average (shown in Figure 7) and use it as input to Algorithm 1 to find the optimal value of m for each team.

More specifically, let the MPDATA domain be of size $240 \times 240 \times 128$. Then, we consider our four teams as four identical abstract processors, $p = 4$, the speed of each of which is given by the speed function shown in Figure 7. Note that in this function, the amount of workload is given in frames of cells of size 120×128 , while the speed is given in cells per second. As pointed in Section 4, despite the unit of workload used to measure the speed (axis y) is 120×128 times greater than the unit of workload used to measure the size of workload (axis x), we can safely use this function as input to Algorithm 1.




Figure 6: Speeds of four teams built simultaneously as functions of parameter m ($n = 120$ and $l = 128$)




Figure 7: Averaged speed of teams built as a function of parameter m ($m = 120$ and $l = 128$)

The solution returned by Algorithm 1 allocates $m = 112$ frames to even abstract processors and $m = 128$ frames to odd processors. This corresponds to partitioning of the $240 \times 240 \times 128$ domain into two sub-domains of size $120 \times 112 \times 128$, allocated to teams T_0 and T_2 , and two sub-domains of size $120 \times 128 \times 128$, allocated to teams T_1 and T_3 . The traditional load-balanced approach partitions the domain in four equal sub-domains of size $120 \times 120 \times 128$. This is illustrated in Figure 8.




Figure 8: Optimal partitioning of MPDATA of size $240 \times 240 \times 128$ between 4 teams

The theoretical execution time of the even partitioning is 1.486 seconds, while the theoretical execution time of the uneven partitioning returned by Algorithm 1 is 1.386:

$$t = \max\left(\frac{x_r}{s(x_r)}, \frac{x_l}{s(x_l)}\right)$$

$$t = \max\left(\frac{120 \cdot 128 \cdot 128}{1418579}, \frac{120 \cdot 112 \cdot 128}{1436742}\right) = 1.386[s].$$

6. EXPERIMENTAL RESULTS

In this section, we experimentally evaluate the optimization technique presented in Section 5.

The performance results presented in this section are obtained for double precision MPDATA computations corresponding to 40 time steps. All the benchmarks are compiled as native executables using the Intel compiler (v.15.0.2), and run on the Intel Xeon Phi 7120P coprocessor. To ensure the reliability of the results, measurements are repeated multiple times, and average execution times are used. We find the confidence interval and stop the measurements if the sample mean lies in the interval with the confidence level 95%. We use Student's t -test, assuming that the individual observations are independent and their population follows the normal distribution.

Table 1 includes both theoretical and experimental execution times of MPDATA for the domain of size $240 \times 240 \times 128$. These results are obtained for different configurations of partitioning, including the traditional "load-balanced" partitioning ($\Delta m = 0$) and a range of "unbalanced" partitioning for different $\Delta m > 0$. The theoretically optimal Δm returned by Algorithm 1 is equal to 8, which corresponds to the configuration where each odd or even team processes the sub-domain of size $120 \times 128 \times 128$ or $120 \times 112 \times 128$ respectively. In this case, the estimated execution time of 1.386 seconds is very close to the real computation time which is 1.364 seconds. According to experiments, the shortest execution time is achieved for $\Delta m = 9$, when computations take 1.348 seconds.

Table 1: Theoretical and experimental execution times for MPDATA domain of size $240 \times 240 \times 128$ with different configurations of partitioning. The odd work teams process the sub-domain of size $120 \times (120 + \Delta m) \times 128$, while the even teams – $120 \times (120 - \Delta m) \times 128$.

Offset Δm	Theoretical time [s]	Experimental time [s]	Speedup
0	1.486	1.548	1.000
4	1.470	1.470	1.053
6	1.401	1.374	1.127
7	1.422	1.361	1.137
8	1.386	1.364	1.135
9	1.398	1.348	1.148
10	1.397	1.352	1.145
11	1.429	1.372	1.129
12	1.402	1.368	1.131

Comparing the experimental and theoretical times, we can see that the accuracy of theoretical prediction is very good, with prediction errors being as small as 2 – 4%. In general, we can identify two main factors contributing into the prediction error:

- While the experimentally built speed functions of teams T_0, T_1, T_2 and T_3 are not identical, suggesting some degree of their heterogeneity in execution of the MPDATA

workload, our theoretical model considers them homogeneous and represents their speed by the average of the real speed functions, which is then used as input to Algorithm 1.

- During the construction of the speed functions, the speed of a team for problem size $n \times m \times l$ is measured when other teams process in parallel sub-domains of the same size, $n \times m \times l$. However, during the execution of the application in our experiments different teams process sub-domains of slightly different sizes when $\Delta m \neq 0$.

Table 1 also demonstrates the performance gain from applying the proposed load-imbalancing optimization. For the imbalanced configurations presented in this table, we notice a better performance than for the load-balanced configuration of the MPDATA decomposition. The largest performance gain is achieved for $\Delta m = 9$, giving the speedup of 1.148x.

Table 2: Experimental time for all work teams with different partitionings: the odd work teams process the sub-domain of size $120 \times (120 + \Delta m) \times 128$, while the even teams $-120 \times (120 - \Delta m) \times 128$.

Offset Δm	Experimental time [s]				
	Team 0	Team 1	Team 2	Team 3	Total
0	1.515	1.498	1.518	1.503	1.548
4	1.456	1.247	1.455	1.249	1.470
6	1.364	1.161	1.359	1.162	1.374
7	1.355	1.161	1.341	1.168	1.361
8	1.355	1.166	1.349	1.172	1.364
9	1.340	1.155	1.335	1.161	1.348
10	1.345	1.141	1.337	1.152	1.352
11	1.363	1.156	1.357	1.154	1.372
12	1.360	1.163	1.353	1.165	1.368

Table 2 complements the results in Table 1 giving experimental execution times of the individual teams. We can clearly see a significant difference between the execution times measured for the odd and even teams when $\Delta m \neq 0$. Obviously, this difference is caused by the unbalanced workloads for the odd and even teams. However, the total execution time is shorter than in the case of balanced workloads ($\Delta m = 0$).

Table 2 also shows that the total execution time is always slightly longer than the maximum time among all teams. It is mainly due to the fact that the computation time of every team is measured without the overheads of inter-team synchronization required after each time step. In addition, the results in Table 2 are presented in Figure 9 in a graphical form.

Finally, we evaluate the proposed model-based partitioning algorithm for the MPDATA domain of size $480 \times 480 \times 128$. As in the previous case, the application is executed for different configurations of partitioning, for a range of Δm . In this case, however, the theoretically optimal configuration returned by Algorithm 1 is exactly the same as the experimentally optimal one, both achieved when $\Delta m = 20$. The




Figure 9: Experimental execution times measured for individual work teams and the total execution time measured for the whole MPDATA workload

prediction errors are also smaller in this case, not exceeding 3%. The experimental execution time for $\Delta m = 20$ is 5.338 seconds, in comparison with 6.140 seconds for the even partitioning. This allows us to accelerate the MPDATA computations by 1.15 times. Moreover, the performance gain is also observed for other unbalanced configurations, but it is smaller than 1.15x. The results of these experiments are included in Table 3.

Table 3: Theoretical and experimental execution times for MPDATA domain of size $480 \times 480 \times 128$ with different configurations of partitioning. The odd work teams process the sub-domain of size $240 \times (240 + \Delta m) \times 128$, while the even teams $-240 \times (240 - \Delta m) \times 128$.

Offset Δm	Theoretical time [s]	Experimental time [s]	Speedup
0	6.136	6.140	1.000
4	5.731	5.681	1.081
8	5.809	5.806	1.058
12	5.543	5.453	1.126
16	5.509	5.418	1.133
20	5.499	5.338	1.150
24	5.624	5.477	1.121

7. CONCLUSION

Modern compute nodes are characterized by both the increasing number of (possibly, heterogeneous) processing elements and a high level of complexity of their integration. Various resources such as caches and data links are shared in an hierarchical and non-uniform way. This makes the development of efficient applications for such platforms a very difficult and challenging task. It would be naive to expect that the performance profile of real-life scientific applications on these platforms will always be comfortably nice and smooth to suit traditional load-balancing techniques used for minimization of their computation time. Therefore, new optimization approaches that do not rely on such increasingly unrealistic assumptions are needed. This work has presented one such approach and demonstrated its applicability to optimization of a real-life application on a modern HPC platform.

Acknowledgments

This research was conducted with the financial support of NCN under grant no. UMO-2011/03/B/ST6/03500. We gratefully acknowledge the help and support provided by Jamie Wilcox from Intel EMEA Technical Marketing HPC Lab. This work is partially supported by EU under the COST Program Action IC1305: Network for Sustainable Ultrascale Computing (NESUS).

8. REFERENCES

- [1] *Parallel Programming and Optimization with Intel Xeon Phi Coprocessors, Handbook on the Development and Optimization of Parallel Applications for Intel Xeon Processors and Intel Xeon Phi Coprocessors*. Colfax International, 2013.
- [2] A. AlOnazi, D. Keyes, A. Lastovetsky, and V. Rychkov. Design and optimization of openfoam-based cfd applications for hybrid and heterogeneous hpc platforms. *arXiv preprint arXiv:1505.07630*, 2015.
- [3] C. Augonnet, S. Thibault, and R. Namyst. Automatic calibration of performance models on heterogeneous multicore architectures. In *EuroPar'09*, pages 56–65, 2009.
- [4] M. Cierniak, M. Zaki, and W. Li. Compile-Time Scheduling Algorithms for Heterogeneous Network of Workstations. *Computer J.*, 40:356–372, 1997.
- [5] D. Clarke, A. Lastovetsky, and V. Rychkov. Dynamic load balancing of parallel computational iterative routines on highly heterogeneous hpc platforms. *Parallel Processing Letters*, 21(02):195–217, 2011.
- [6] D. Clarke, A. Lastovetsky, and V. Rychkov. Column-based matrix partitioning for parallel matrix multiplication on heterogeneous processors based on functional performance models. In *Euro-Par 2011: Parallel Processing Workshops*, pages 450–459. Springer, 2012.
- [7] J. Colaco, A. Matoga, et al. Transparent application acceleration by intelligent scheduling of shared library calls on heterogeneous systems. In *PPAM 2013, Part I*, pages 693–703, 2014.
- [8] M. Fatica. Accelerating Linpack with CUDA on heterogeneous clusters. In *GPGPU-2*, pages 46–51. ACM, 2009.
- [9] A. Ilic, F. Pratas, P. Trancoso, and L. Sousa. High-performance computing on heterogeneous systems: Database queries on cpu and gpu. In *High Performance Scientific Computing with Special Emphasis on Current Capabilities and Future Perspectives*. IOS Press, 2011.
- [10] A. Kalinov and A. Lastovetsky. Heterogeneous distribution of computations solving linear algebra problems on networks of heterogeneous computers. *Journal of Parallel and Distributed Computing*, 61(4):520–535, 2001.
- [11] A. Lastovetsky and R. Reddy. Data partitioning with a realistic performance model of networks of heterogeneous computers. In *Proceedings of the 18th International Parallel and Distributed Processing Symposium (IPDPS 2004)*, Santa Fe, New Mexico, USA, 26-30 April 2004. IEEE Computer Society.
- [12] A. Lastovetsky and R. Reddy. Data distribution for dense factorization on computers with memory heterogeneity. *Parallel Computing*, 33(12):757–779, 2007.
- [13] A. Lastovetsky and R. Reddy. Data partitioning with a functional performance model of heterogeneous processors. *International Journal of High Performance Computing Applications*, 21(1):76–90, 2007.
- [14] A. Lastovetsky and J. Twamley. Towards a realistic performance model for networks of heterogeneous computers. In *High Performance Computational Science and Engineering*, pages 39–57. Springer, 2005.
- [15] M. D. Linderman, J. D. Collins, H. Wang, et al. Merge: a programming model for heterogeneous multi-core systems. *SIGPLAN Not.*, 43:287–296, 2008.
- [16] C.-K. Luk, S. Hong, and H. Kim. Qilin: Exploiting parallelism on heterogeneous multiprocessors with adaptive mapping. In *MICRO-42*, pages 45–55, 2009.
- [17] J. Martínez, E. Garzón, A. Plaza, and I. García. Automatic tuning of iterative computation on heterogeneous multiprocessors with ADITHE. *J. Supercomput.*, 58(2):151–159, 2011.
- [18] Y. Ogata, T. Endo, N. Maruyama, and S. Matsuoka. An efficient, model-based CPU-GPU heterogeneous FFT library. In *IPDPS 2008*, pages 1–10, 2008.
- [19] Z. Piotrowski, A. Wyszogrodzki, and P. Smolarkiewicz. Towards Petascale Simulation of Atmospheric Circulations with Soundproof Equations. *Acta Geophysica*, 59:1294–1311, 2011.
- [20] G. Quintana-Ortí, F. D. Igual, E. S. Quintana-Ortí, and R. A. van de Geijn. Solving dense linear systems on platforms with multiple hardware accelerators. *SIGPLAN Not.*, 44:121–130, 2009.
- [21] P. Smolarkiewicz. Multidimensional Positive Definite Advection Transport Algorithm: An Overview. *Int. J. Numer. Meth. Fluids*, 50:1123–1144, 2006.
- [22] P. Smolarkiewicz and W. Grabowski. The Multidimensional Positive Definite Advection Transport Algorithm: Nonoscillatory Option. *J. Comput. Phys.*, 86:355–375, 1990.
- [23] P. Smolarkiewicz and L. Margolin. MPDATA-A Multipass Donor Cell Solver for Geophysical Flows. In E. Toro, editor, *Godunov Methods*, pages 833–839. Springer US, 2001.
- [24] F. Song, S. Tomov, and J. Dongarra. Enabling and scaling matrix computations on heterogeneous multi-core and multi-GPU systems. In *ICS '12*, pages 365–376. ACM, 2012.
- [25] L. Szustak, K. Rojek, and P. Gepner. Using Intel Xeon Phi coprocessor to accelerate computations in MPDATA algorithm. *Lect. Notes in Comp. Sci.*, 8385:582–592, 2014.
- [26] L. Szustak, K. Rojek, T. Olas, L. Kuczynski, K. Halbiniak, and P. Gepner. Adaptation of MPDATA heterogeneous stencil computation to Intel Xeon Phi coprocessor. *Scientific Programming*, In Print, 2015.
- [27] C. Yang, F. Wang, Y. Du, et al. Adaptive optimization for petascale heterogeneous CPU/GPU computing. In *Cluster'10*, pages 19–28, 2010.
- [28] Z. Zhong, V. Rychkov, and A. Lastovetsky. Data partitioning on multicore and multi-gpu platforms using functional performance models. *IEEE*

Transactions on Computers, pages 1–14, 12/2014.

- [29] Z. Zhong, V. Rychkov, and A. Lastovetsky. Data partitioning on heterogeneous multicore platforms. In *2011 IEEE International Conference on Cluster Computing (Cluster 2011)*, pages 580–584, Austin, Texas, USA, Sept 26-30 2011. IEEE Computer Society.